

CLAIMS:

1. A method for breathing of scheduling algorithms for a storage device, the method comprising:
 - (a) computing a worst-case duration of a breathing cycle for the storage device, the worst-case duration being referred to by P;
 - (b) starting a breathing cycle;
 - (c) determining if one of the following becomes true before the end of P time units:
 - (i) a number of real-time requests is at least a predetermined threshold based on a number of data streams and performance parameters of the storage device; and
 - (ii) a number of pending requests for any single stream becomes more than one;
 - (d) if at least one of (i) and (ii) remain true during the duration of P time units from the start of the breathing cycle, starting a subsequent breathing cycle after completion of the breathing cycle; and
 - (e) if both of (i) and (ii) are not true during the duration of P time units from the start of the breathing cycle, waiting P time units from the start of the breathing cycle before starting the subsequent breathing cycle.
2. The method of claim 1, wherein best-effort requests arriving during the breathing cycle are handled during the breathing cycle.
3. The method of claim 1, further comprising, repeating steps (b) - (e) for a plurality of breathing cycles.
4. The method of claim 1, further comprising (f) calculating an actual bit-rate of a data stream based on the determination of step (ii).
5. The method of claim 4, further comprising (g) changing a bit rate for the data stream based on the calculating in step (f).
6. The method of claim 5, wherein step (g) comprises reserving a higher bit rate for the data stream where its estimated maximum bit rate is exceeded.

7. The method of claim 6, wherein the data stream is transferred between the storage device and a buffer, the method further comprises (h) increasing a size of the buffer.
8. The method of claim 5, wherein step (g) comprises reserving a lower bit rate for the data stream where its estimated maximum bit rate is not exceeded.
9. The method of claim 8, wherein the data stream is transferred between the storage device and a buffer, the method further comprises (h) decreasing a size of the buffer.
10. A storage device scheduler for controlling the breathing of scheduling algorithms for a storage device, the storage device scheduler comprising:
 - (a) means for computing a worst-case duration of a breathing cycle for the storage device, the worst-case duration being referred to by P;
 - (b) means for instructing the starting of a breathing cycle;
 - (c) means for determining if one of the following becomes true before the end of P time units:
 - (i) a number of real-time requests is at least a predetermined threshold based on a number of data streams and performance parameters of the storage device; and
 - (ii) a number of pending requests for any single stream becomes more than one;
 - (d) means for starting a subsequent breathing cycle after completion of the breathing cycle if at least one of (i) and (ii) remain true during the duration of P time units from the start of the breathing cycle;; and
 - (e) means for waiting P time units from the start of the breathing cycle before starting the subsequent breathing cycle if both of (i) and (ii) are not true during the duration of P time units from the start of the breathing cycle.
11. The storage device scheduler of claim 10, further comprising, repeating steps (b) - (e) for a plurality of breathing cycles.
12. The storage device scheduler of claim 10, further comprising (f) calculating an actual bit-rate of a data stream based on the determination of step (ii).

13. The storage device scheduler of claim 12, further comprising (g) changing a bit rate for the data stream based on the calculating in step (f).

14. The storage device scheduler of claim 13, wherein step (g) comprises reserving a higher bit rate for the data stream where its estimated maximum bit rate is exceeded.

15. The storage device scheduler of claim 14, wherein the data stream is transferred between the storage device and a buffer, the method further comprises (h) increasing a size of the buffer.

16. The storage device scheduler of claim 13, wherein step (g) comprises reserving a lower bit rate for the data stream where its estimated maximum bit rate is not exceeded.

17. The storage device scheduler of claim 16, wherein the data stream is transferred between the storage device and a buffer, the method further comprises (h) decreasing a size of the buffer.

18. The storage device scheduler of claim 10, further comprising means for handling best-effort requests arriving during the breathing cycle during the breathing cycle.

19. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for breathing of scheduling algorithms for a storage device, the method comprising:

(a) computing a worst-case duration of a breathing cycle for the storage device, the worst-case duration being referred to by P;

(b) starting a breathing cycle;

(c) determining if one of the following becomes true before the end of P time units:

(i) a number of real-time requests is at least a predetermined threshold based on a number of data streams and performance parameters of the storage device; and

(ii) a number of pending requests for any single stream becomes more than one;

(d) if at least one of (i) and (ii) remain true during the duration of P time units from the start of the breathing cycle, starting a subsequent breathing cycle after completion of the breathing cycle; and

(e) if both of (i) and (ii) are not true during the duration of P time units from the start of the breathing cycle, waiting P time units from the start of the breathing cycle before starting the subsequent breathing cycle.

20. A computer program product embodied in a computer-readable medium for breathing of scheduling algorithms for a storage device, the method comprising:

(a) computer readable program code means for computing a worst-case duration of a breathing cycle for the storage device, the worst-case duration being referred to by P;

(b) computer readable program code means for starting a breathing cycle;

(c) computer readable program code means for determining if one of the following becomes true before the end of P time units:

(i) a number of real-time requests is at least a predetermined threshold based on a number of data streams and performance parameters of the storage device; and

(ii) a number of pending requests for any single stream becomes more than one;

(d) computer readable program code means for if at least one of (i) and (ii) remain true during the duration of P time units from the start of the breathing cycle, starting a subsequent breathing cycle after completion of the breathing cycle; and

(e) computer readable program code means for if both of (i) and (ii) are not true during the duration of P time units from the start of the breathing cycle, waiting P time units from the start of the breathing cycle before starting the subsequent breathing cycle.